Various Approaches to Data Association

Peter Willett

University of Connecticut

NATO SET-157
Multisensor Fusion:
Advanced Methodology and Applications

Outline

1. Data Association
   - GNN
   - PDAF
   - MHT
   - MFA
2. New Trackers
   - PHD
   - MLPDA
3. Results
4. Label-Free Estimation
   - MMOSPA Estimation
5. Data Association Without Labels
   - SJPDA
   - Labeling a Label-Free Tracker
   - SMHT

What Is Data Association?



- Which measurements come from which targets?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD
- (How many targets are there?)

What Is Data Association?



- Which measurements come from which targets?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD
- (How many targets are there?)

What Is Data Association?



- Which measurements come from which targets?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD
- (How many targets are there?)

What Is Data Association?



- Which measurements come from which targets?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD
- (How many targets are there?)

What Is Data Association?



- Which measurements come from which targets?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD
- (How many targets are there?)

What Is Data Association?



- Which measurements come from which targets?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD
- (How many targets are there?)

What Is Data Association?



- Which measurements come from which targets?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD
- (How many targets are there?)

What Is Data Association?



- Which measurements come from which targets?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD
- (How many targets are there?)

What Is Data Association?



- Which measurements come from which targets?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD
- (How many targets are there?)

What Is Data Association?



- Which measurements come from which targets?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD
- (How many targets are there?)

What Is Data Association?



- Which measurements come from which targets?
- Commonly accepted methods:
  - (J)PDAF
  - GNN
  - MHT (HO-MHT or TO-MHT)
  - MFA
- Emerging methods:
  - PMHT
  - PHD
  - MLPDA
  - JMPD
- (How many targets are there?)

What Isn't Data Association?



- linear or nonlinear dynamics and/or measurement model, and added measurement noise can be Gaussian or not
- if both linear and Gaussian, then it's time for the Kalman filter
- data association: special non-Gaussian measurement noise
- extended Kalman filters (EKFs), unscented Kalman filters (UKFs) and particle filters (PFs) are general
- but there is no "magic bullet"
- no naked EKF, UKF or PF works with data association

What Isn't Data Association?



- linear or nonlinear dynamics and/or measurement model, and added measurement noise can be Gaussian or not
- if both linear and Gaussian, then it's time for the Kalman filter
- data association: special non-Gaussian measurement noise
- extended Kalman filters (EKFs), unscented Kalman filters (UKFs) and particle filters (PFs) are general
- but there is no "magic bullet"
- no naked EKF, UKF or PF works with data association

What Isn't Data Association?



- linear or nonlinear dynamics and/or measurement model, and added measurement noise can be Gaussian or not
- if both linear and Gaussian, then it's time for the Kalman filter
- data association: special non-Gaussian measurement noise
- extended Kalman filters (EKFs), unscented Kalman filters (UKFs) and particle filters (PFs) are general
- but there is no "magic bullet"
- no naked EKF, UKF or PF works with data association

What Isn't Data Association?



- linear or nonlinear dynamics and/or measurement model, and added measurement noise can be Gaussian or not
- if both linear and Gaussian, then it's time for the Kalman filter
- data association: special non-Gaussian measurement noise
- extended Kalman filters (EKFs), unscented Kalman filters (UKFs) and particle filters (PFs) are general
- but there is no "magic bullet"
- no naked EKF, UKF or PF works with data association

What Isn't Data Association?



- linear or nonlinear dynamics and/or measurement model, and added measurement noise can be Gaussian or not
- if both linear and Gaussian, then it's time for the Kalman filter
- data association: special non-Gaussian measurement noise
- extended Kalman filters (EKFs), unscented Kalman filters (UKFs) and particle filters (PFs) are general
- but there is no "magic bullet"
- no naked EKF, UKF or PF works with data association

What Isn't Data Association?



- linear or nonlinear dynamics and/or measurement model, and added measurement noise can be Gaussian or not
- if both linear and Gaussian, then it's time for the Kalman filter
- data association: special non-Gaussian measurement noise
- extended Kalman filters (EKFs), unscented Kalman filters (UKFs) and particle filters (PFs) are general
- but there is no "magic bullet"
- no naked EKF, UKF or PF works with data association

Data Association Rules

- at most one measurement per scan per target (per sensor)
- target detected with probability P_d
- number of false alarms Poisson with mean λV
- false alarms are each usually distributed uniformly, but non-homogeneous clutter is fine
- PMHT measurement model is different but convenient, not widely accepted

Data Association Rules

- at most one measurement per scan per target (per sensor)
- target detected with probability P_d
- number of false alarms Poisson with mean λV
- false alarms are each usually distributed uniformly, but non-homogeneous clutter is fine
- PMHT measurement model is different but convenient, not widely accepted

Why is Data Association Hard?

Even with perfect detection (P_d=1)

$p(x_k | z^k)$

Why is Data Association Hard?

Even with perfect detection (P_d=1) . . .



Why is Data Association Hard?

Even with perfect detection (P_d=1) . . .



Why is Data Association Hard?

Even with perfect detection (P_d=1) . . .



The point is that optimal posterior estimation is doomed due to the "explosion" of modes.

Global Nearest Neighbor (GNN)



NNF: Minimize cost for each filter.

GNN: Minimize global cost.

Global Nearest Neighbor (GNN)



NNF: Minimize cost for each filter.

GNN: Minimize global cost.

Global Nearest Neighbor (GNN)



NNF: Minimize cost for each filter.

GNN: Minimize global cost.

Global Nearest Neighbor (GNN)



NNF: Minimize cost for each filter.

GNN: Minimize global cost.

Global Nearest Neighbor (GNN)



NNF: Minimize cost for each filter.

GNN: Minimize global cost.

Global Nearest Neighbor (GNN)



NNF: Minimize cost for each filter.

GNN: Minimize global cost.

The Probabilistic Data Association Filter



Joint (J)PDAF: for multiple targets, weights are computed from joint association events, and tracks are updated individually.

The PDAF is perhaps the only consistent tracker.

Hypothesis-Oriented Multi-Hypothesis Tracker (HO-MHT)



Track-Oriented Multi-Hypothesis Tracker (TO-MHT)



T_i = target I
z_j = measurement j
ϕ = missed detection

Track-Oriented Multi-Hypothesis Tracker (TO-MHT)

Kalman & likelihood (P_d & λ) via Reid



only one of these can be true

T_i = target I
z_j = measurement j
ϕ = missed detection

Track-Oriented Multi-Hypothesis Tracker (TO-MHT)

Kalman & likelihood (P_d & λ) via Reid



only one of these can be true

T_i = target I
z_j = measurement j
ϕ = missed detection

Track-Oriented Multi-Hypothesis Tracker (TO-MHT)

Kalman & likelihood (P_d & λ) via Reid



only one of these can be true

T_i = target I
z_j = measurement j
ϕ = missed detection

Multiple-Frame Assignment (MFA)



- maximize LLR with constraint that no more than one contact per target per scan
- subgradient
- "scan" can be a sensor
- MFA = TO-MHT

Multiple-Frame Assignment (MFA)



- maximize LLR with constraint that no more than one contact per target per scan
- subgradient
- "scan" can be a sensor
- MFA = TO-MHT

Multiple-Frame Assignment (MFA)



- maximize LLR with constraint that no more than one contact per target per scan
- subgradient
- "scan" can be a sensor
- MFA = TO-MHT

MFA Optimization

Just to be mathematical for a slide, the problem is the *constrained* optimization

$$\min_{\{\rho\}} \sum_{i_1=0}^{n_1} \sum_{i_2=0}^{n_2} \cdots \sum_{i_S=0}^{n_S} c_{i_1 i_2 \ldots i_S} \rho_{i_1 i_2 \ldots i_S}$$

In this expression $c_{i_1 i_2 \ldots i_S}$ is the cost of each association, given by the negative log-likelihood ratio (NLLR)[1] In fact, this last really amounts to the likelihood function of a set of Kalman filters – one for each association of track i_1 to $z_{i_2}(k), z_{i_3}(k+1), \ldots, z_{i_S}(k+S-1)$ ($i_j = 0$ means a missed detection). It is estimated that more than 90% of the effort of the optimization is the calculation of these filters.

[1] It is a *ratio* since it is conveniently divided by the likelihood of all measurements being false alarms.

The Constraint

$$\sum_{i_2=0}^{n_2}\sum_{i_3=0}^{n_3}\cdots\sum_{i_S=0}^{n_S}\rho_{i_1 i_2\ldots i_S}=1 \qquad i_1=1,2,\ldots,n_1$$

$$\sum_{i_1=0}^{n_1}\sum_{i_3=0}^{n_3}\cdots\sum_{i_S=0}^{n_S}\rho_{i_1 i_2\ldots i_S}=1 \qquad i_2=1,2,\ldots,n_2$$

$$\vdots$$

$$\sum_{i_1=0}^{n_1}\sum_{i_2=0}^{n_2}\cdots\sum_{i_{S-1}=0}^{n_{S-1}}\rho_{i_1 i_2\ldots i_S}=1 \qquad i_S=1,2,\ldots,n_S$$

- For $S=2$ a valid ρ might be as on the right.
- But for $S=2$ this is GNN, and the solution is "easy".
- Auction or JVC: polynomial time.

miss — measurements

false alarm — targets

	0	1	2	3	4	5	6
0		1	1	0	0	1	0
1	0	0	0	0	1	0	0
2	1	0	0	0	0	0	0
3	0	0	0	1	0	0	0
4	1	0	0	0	0	0	0
5	0	0	0	0	0	0	1

The Probability Hypothesis Density (PHD)



Figure: Prediction

- Most trackers ask: Where are targets 1, 2 and 3?
- PHD asks: Is there a target in this location?
- PHD integrated over a volume gives expected number of targets therein.
- We think of this as predicting and updating the probabilities that infinitesimal “cans” contain a target.

The Probability Hypothesis Density (PHD)



Figure: Update

- Measurements provide Bayesian update of the probabilities.
- Birth, death and spawn can each be accounted for.
- As can volume goes to zero, it converges to Mahler's PHD.

GM-CPHD

- Target Death Problem: targets die faster than they should when a detection is missed.
- Cardinalized (CPHD) introduced HMM on overall target number, which addresses target-death problem.
- "Spooky Action at a Distance" (Ulmke et al.), a phenomenon that we are coming more and more to respect and to fear.
- GM-PHD surface can be approximated by Gaussian Mixture as opposed to particle update.

GM-CPHD

- Target Death Problem: targets die faster than they should when a detection is missed.
- Cardinalized (CPHD) introduced HMM on overall target number, which addresses target-death problem.
- "Spooky Action at a Distance" (Ulmke et al.), a phenomenon that we are coming more and more to respect and to fear.
- GM-PHD surface can be approximated by Gaussian Mixture as opposed to particle update.

GM-CPHD

- Target Death Problem: targets die faster than they should when a detection is missed.
- Cardinalized (CPHD) introduced HMM on overall target number, which addresses target-death problem.
- “Spooky Action at a Distance” (Ulmke et al.), a phenomenon that we are coming more and more to respect and to fear.
- GM-PHD surface can be approximated by Gaussian Mixture as opposed to particle update.

GM-CPHD

- Target Death Problem: targets die faster than they should when a detection is missed.
- Cardinalized (CPHD) introduced HMM on overall target number, which addresses target-death problem.
- “Spooky Action at a Distance” (Ulmke et al.), a phenomenon that we are coming more and more to respect and to fear.
- GM-PHD surface can be approximated by Gaussian Mixture as opposed to particle update.

MLPDA Background

MLPDA estimator/detector was first proposed by Jauffret and Bar-Shalom in 1990 as an approach for passive narrowband target motion analysis. Subsequent work has expanded it to tracking applications with radar and active sonar data.

The idea behind MLPDA is actually very simple. We assume we have a target with some parameterized, deterministic motion in the presence of clutter. With some basic assumptions, we can formulate a likelihood function and find the maximum.

MLPDA Assumptions

- Single target present in each frame with known detection probability P_d. Detections are independent across frames.
- Zero or one measurement per frame from target.
- Kinematics of target are deterministic. Usually parameterized as a straight line, although any other parameterization is valid.
- False detections are uniformly distributed in search volume.
- Poisson (λ) number of false detections.
- Target measurements have additive zero-mean Gaussian noise.
- Measurements at different times, conditioned on the parameterized state, are independent.

MLPDA Log-Likelihood Function

Find the parameterized state X that maximizes the LLR:

$$\Lambda(Z, a|X) = \sum_{i=1}^{N_w} \ln \left[1 + \frac{P_d}{\lambda(1 - P_d)} \sum_{j=1}^{m_i} \rho_j(i) p(z_j(i)|X) \right]$$

where

- N_w – number of frames in a window
- m_i – number of measurements in the ith frame
- $\rho_j(i)$ – amplitude likelihood ratio for the jth measurement in the ith frame.
- λ – spatial clutter density
- P_d – probability of detection
- $p(z_j(i)|X)$ – target-centered Gaussian, conditioned on (the unknown) initial target state parameter vector X

Metron Dataset

- Simulated sonar data
- Ping period = 180 sec
  Number of pings = 200
- TDOA error: 0.4 sec
  Bearing error: 8°
  Doppler error: 0.5 m/sec
- Sensor grid
  4 Sources (FM, CW)
  25 Receivers
- 5 Scenarios
- Number false alarms per receiver per scan ≈ 20
- $P_d \approx$ 15%

	Target1	Target2
Speed	6	6
Start	(36000, 49000)	(48000, 49000)
Box	13500 x 13500	13500 x 13500
Motion	Clockwise	CounterClockwise
Label	0	1

	Target3	Target4
Speed	3	3
Start	(24250, 30250)	(47750, 30250)
Box	6750 x 6750	6750 x 6750
Motion	CounterClockwise	Clockwise
Label	1	0

Table: Ground Truth in Scenario1.
Red denotes a surprise.

Metron Scenario 1 – all measurements in one window of data



Scenario 1



Figure: Scenario 1.

GM-CPHD Results – Scenario 1



Figure: Scans 1-50



Figure: Scans 50-100

GM-CPHD Results – Scenario 1



Figure: Scans 100-150



Figure: Scans 150-200

GM-CPHD Results – Scenarios 2, 3



Figure: Scenario 2, three targets (two inner rectangles and one outer).



Figure: Scenario 3, a bow-tie and a target crossing at an oblique angle.

GM-CPHD Movie – Scenario 4

MLPDA Results – Scenario 1



Figure: Scenario 1

MLPDA Results – Scenarios 2, 3



Figure: Scenario 2



Figure: Scenario 3

Rockets & torpedoes

There are some circumstances in which multi-object estimation ought to be “label-free”. One wants to track all objects, but one is not especially concerned which one is which.

MMSE Estimation

Assuming that MMSE estimation is the goal

$$\{\hat{x}_i\}_{i=1}^n = \arg\min \left\{ \int \sum_{i=1}^n |x_i - \hat{x}_i|^2 p(x_1, \ldots, x_n) dx_1 \ldots dx_n \right\}$$

This will be a shorthand notation: in fact, for estimation, we would really have

$$\begin{aligned} &\{\hat{x}_i(z)\}_{i=1}^n \\ &= \arg\min \left\{ \int \sum_{i=1}^n |x_i - \hat{x}_i(z)|^2 p(x_1, \ldots, x_n | z) dx_1 \ldots dx_n \right\} \end{aligned}$$

in which z is an observation – that is, we are doing a-posteriori estimation.

Constrained MMSE Estimation

One does not usually write this, but when one estimates in the multi-object situation, one usually has the “labeling” constraints

$$x_1 \sim p_1, \qquad x_2 \sim p_2, \qquad \ldots \qquad x_n \sim p_n$$

as opposed to, say

$$x_3 \sim p_1, \qquad x_n \sim p_2, \qquad \ldots \qquad x_1 \sim p_n$$

Recalling that this is a minimization problem, the constraints can be doing no good as regards our minimization criterion (MSE). Are they doing any harm?



What is wrong with MMSE?

The problem, of course, is that the MMSE criterion

$$\{\hat{x}_i\}_{i=1}^n = \arg\min \left\{ \int \sum_{i=1}^n |x_i - \hat{x}_i|^2 \, p(x_1, \ldots, x_n) dx_1 \ldots dx_n \right\}$$

has little meaning *without* labeling. What can one do?

It is clear that RFS will arise. But we should first note that Drummond[2] discussed measures that do not overly penalize labeling mistakes by calculating the SE of the best assignment.

[2] Drummond & Fridling, "Ambiguities in evaluating performance of multiple target tracking algorithms," *SPIE Conf. on Signal and Data Processing of Small Targets*, 1992.

The OSPA Metric

The Optimal Mass Transfer (OMAT[3]) *metric* is less ad-hoc than assignment, but there are weaknesses. Here we refer to Optimal Subpattern Assignment (OSPA[4]) that is more general and encompassing.

$$\tilde{d}_p^{(c)}(\widehat{\mathbf{X}}, \mathbf{X}) = \left(\frac{1}{n} \left(\min_{\pi \in \Pi_n} \sum_{i=1}^{n} d^{(c)}(\mathbf{x}^{(i)}, \hat{\mathbf{x}}^{\pi(i)})^p \right) \right)^{1/p}$$

Here, $d^{(c)}(\mathbf{x}, \hat{\mathbf{x}}) \triangleq \min(c, d(\mathbf{x}, \hat{\mathbf{x}}))$ is the distance d between $\mathbf{x}$ and $\hat{\mathbf{x}}$, cut-off at c and $\hat{\mathbf{x}}^{\pi(i)}$ describes the ith permutation (target-reordered) of the vector $\hat{\mathbf{x}}$.

[3] Hoffman & Mahler, "Multitarget Miss Distance via Optimal Assignment," *IEEE T-SMC:A*, May, 2000.

[4] Schuhmacher, Vo & Vo, "A Consistent Metric for Performance Evaluation of Multi-Object Filters," *IEEE T-SP*, 2008.

MMOSPA

- Since one can calculate the squared error (SE), and then minimize its mean value (MSE) to derive the MMSE, in obvious parallel here we are interested in minimizing the *mean* OSPA (MOSPA) measure (the MMOSPA estimate):

$$\text{MOSPA}_p^{(c)}(\widehat{\mathbf{X}}, p(\mathbf{X})) \triangleq \mathbb{E}_{p(\mathbf{X})}\{\tilde{d}_p^{(c)}\}.$$

- Nice theory about the combination of MMOSPA and RFS[5] and nice algorithms to find the MMOSPA estimates[6] too.

[5] M. Guerriero, D. Svensson, L. Svensson and P. Willett, "Shooting Two Birds with Two Bullets: How to Find Minimum Mean OSPA," FUSION 2010.

[6] D. Crouse, P. Willett, M. Guerriero and L. Svensson, "An Approximate Miminum MOSPA Estimator," ICASSP 2011.

Since this isn't going to be easy ...

Applications where MOSPA is reasonable

- two threats are attacking
- radar cueing: there is no interest in which target is which, but where there are targets is very important
- collision avoidance: it is not of interest to know which car is which, the goal is to avoid all cars
- maritime surveillance: send helicopters as close as possible to a set of vessels with anomalous behavior

But it will turn out that MMOSPA has application in a wider context. Even in situations where target labeling is very much desired.

Example

Two Gaussian-prior targets, ballistic trajectory, separated by unit distance. We assume $P_d = 1$ – two unlabeled hits – and Gaussian observation noise at each scan, of standard deviation σ_w:



Does the Admiral care about MOSPA?

Probably not yet.
So let us define the kill probability as the probability that the estimates of the target locations are within a ball of radius ϵ centered around the true target locations

$$P_{kill}(\hat{\mathbf{x}}, \epsilon) = P(\|\hat{\mathbf{x}} - \mathbf{x}\| \leq \epsilon)$$



The Problem with Closely-Spaced Objects



Top: two targets being tracked by the JPDAF – and believe it or not the *coupled* JPDAF is worse!

Below: the label-free SJPDAF.

As discussed in some detail[a], close tracks tend to merge with soft-association trackers and to repel in hard-association ones.

[a] Coraluppi, Carthel, Willett, Dingboe, O'Neill & Luginbuhl, "The Track Repulsion Effect in Automatic Tracking," *FUSION 2009*.

Soft Assignment Trackers (PDA, JPDA) Merge CSOs

- MMSE target state update
- It is fairly well-known[7] that soft association trackers tend to "coalesce" (merge) nearby targets.



Two targets' pdf contours are notionally represented in red & green. For the "true posterior," the thickness of the lines indicates the posterior weight in the Gaussian-mixture mode. The MMSE estimates are the soft-association trackers' updated pdfs – note that due to the association uncertainty they have tended toward one another.

[7] R. Fitzgerald, "Track Biases and Coalescence with Probabilistic Data Association," IEEE T-AES, 1985.

Hard-Association Trackers (GNN, MHT) Repel CSOs!

- Choose only the association that fits best
- It is less well-known[8] that with hard association trackers nearby targets have difficulty crossing.
- Below: right happens more often than left.





[8] S. Coraluppi, C. Carthel, P. Willett, M. Dingboe, O. O'Neill, and T. Luginbuhl, "The Track Repulsion Effect in Automatic Tracking," *FUSION 2009*.

The SJPDA

- With a multimodal pdf the MMSEE "hedges its bets" and produces an estimate somewhere among the modes
  - This is track coalescence.
- MMOSPA estimation eliminates some of the extra pdf modes, the ones that correspond to re-orderings of the items being estimated
- Since its genesis was in random finite set (RFS) theory, this was given the monicker *Set Joint Probabilistic Data Association*, or SJPDA

The SJPDA and MMOSPA

In the SJPDA context (Gaussian priors, Gaussian measurement noise, measurement-origin uncertainty) the MOSPA error is proportional to the trace of an *unordered* joint covariance matrix,

$$\begin{aligned}
\mathbf{P}(k) &\equiv \mathrm{E}\left[\min_{\mathbf{a}} \left(\mathbf{x}_{\mathbf{a}}(k) - \hat{\mathbf{x}}(k|k)\right)\left(\mathbf{x}_{\mathbf{a}}(k) - \hat{\mathbf{x}}(k|k)\right)^T \middle| Z^k\right] \\
&= \sum_{i=1}^{N_H} \Pr\{\theta_i\}\, \mathrm{E}\left[\min_{\mathbf{a}} (\mathbf{x}_{\mathbf{a}}(k)-\hat{\mathbf{x}}(k|k))(\mathbf{x}_{\mathbf{a}}(k)-\hat{\mathbf{x}}(k|k))^T \middle| \theta_i, Z^k\right] \\
&\approx \sum_{i=1}^{N_H} \Pr\{\theta_i\} \min_{\mathbf{a}_i} \mathrm{E}\left[\left(\mathbf{x}_{\mathbf{a}_i}(k)-\hat{\mathbf{x}}(k|k)\right)\left(\mathbf{x}_{\mathbf{a}_i}(k)-\hat{\mathbf{x}}(k|k)\right)^T \middle| \theta_i, Z^k\right]
\end{aligned}$$

The SJPDAF uses the a_i's chosen to minimize this (approximation's) trace, and then updates the covariance in

$$\sum_{i=1}^{N_H} \Pr\{\theta_i\} \left(\mathbf{P}_{\mathbf{a}_i}(k|k) + (\hat{\mathbf{x}}_{i,\mathbf{a}_i}(k|k) - \hat{\mathbf{x}}(k|k))(\hat{\mathbf{x}}_{i,\mathbf{a}_i}(k|k) - \hat{\mathbf{x}}(k|k))^T\right)$$

The SJPDA Intuition



setup Gaussian prior, 2 targets, $P_d = 0.7$, 2 measurements

top left true posterior

top right MMOSPA posterior

bottom left JPDA posterior

bottom right SJPDA posterior

Nearest Neighbor Approaches For Close Targets

- The ENNPDA[9] or BHTT[10] are ways to get the best of both.
  1. compute all joint likelihoods, as for a GNN.
  2. normalize them to probabilities, as for a JPDA.
  3. select the most probable and discard other associations
  4. PDA update each target based on detection/miss events
- Fitzgerald used "Cheap JPDA", Drummond JVC

[9] R. Fitzgerald, "Development of practical PDA logic for multitarget tracking by microprocessor," *in Multitarget-Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom, Ed., 1990.

[10] O. Drummond, "Best hypothesis target tracking and sensor fusion," *SPIE*, 1999.

The JPDA*

- The ENNPDA to some extent reduces coalescence (w.r.t. JPDA) and reduces track loss (w.r.t. GNN).
- But better performance is possible by less Draconian event-pruning[11].
- Define
  - ϕ is event denoting which targets are detected
  - ψ is event denoting which plots are from targets
  - $A(\phi, \psi)$ is event of associations consistent with ϕ and ψ.
  - Λ is likelihood
- Now consider doing JPDA with only those events $\{A^*\}$:

$$A^*(\phi, \psi) = \arg\max_A \{\Lambda(A(\phi, \psi)\}$$

[11] H. Blom and E. Bloem, "Probabilistic Data Association Avoiding Track Coalescence," *IEEE TAC*, 2000.

Why talk about JPDA*?

- JPDA* is not a label-free tracker.
- There is a related revisit[12] of it that splits events into permutation-variant and permutation-invariant – watch for further development of this.
- The SJPDA* is a hybrid of the SJPDA and JPDA*:
  - do the "most-likely" hypothesis pruning based on the JPDA*
  - after you are done, use the MMOSPA criterion to merge the target-permuted hypotheses, like the SJPDA
- And there are *coupled* versions (like CJPDA) of both of these: CSJPDA & CSJPDA*.
- And I wanted to mention it because it is undeservedly obscure and is pretty neat.

[12] H. Blom and E. Bloem, "Permutation invariance in Bayesian estimation of two targets that maneuver in and out formation flight," FUSION 2009.

Why talk about JPDA*?

- JPDA* is not a label-free tracker.
- There is a related revisit[12] of it that splits events into permutation-variant and permutation-invariant – watch for further development of this.
- The SJPDA* is a hybrid of the SJPDA and JPDA*:
  - do the "most-likely" hypothesis pruning based on the JPDA*
  - after you are done, use the MMOSPA criterion to merge the target-permuted hypotheses, like the SJPDA
- And there are *coupled* versions (like CJPDA) of both of these: CSJPDA & CSJPDA*.
- And I wanted to mention it because it is undeservedly obscure and is pretty neat.

[12] H. Blom and E. Bloem, "Permutation invariance in Bayesian estimation of two targets that maneuver in and out formation flight," FUSION 2009.

But the Admiral Wants Labels ...

An operator might want to:

- Know how many targets there are.
- Know the states of the targets.
- Know what the targets are. This could mean knowing:
  - Their identity relative to a specific time designated by the operator.
  - Their physical nature according to classification (feature) information.
- Have a probabilistic model for the uncertainties in the above estimates.
  - Localization accuracy (covariance)
  - Identity accuracy

The SJPDA With Identity Information

In a particle filter, each particle might represent a probability and a stacked set of states for each target.

$$p_i = \{w_i, \mathbf{x}_i\}$$

For this exposition, consider that a "particle" can be the same as a Gaussian mode in the SJPDAF.
Suppose that we want to be able to change the ordering of the states for the targets without losing or corrupting any information. We can do this by adding an extra component:

$$\tilde{p}_i = \{w_i, \mathbf{x}_i, \mathbf{o}_i\}$$

The vector $\mathbf{o}$ can be an $N_T! \times 1$ vector that lists the probability of a particular ordering of states, initially all zeros and a single 1.

MMOSPA Estimation + Identity

- Suppose that we arrange the *ordering* of the states in a set of particles to coincide with the ordering for MMOSPA estimation.
- If we take the average over the states, we get the MMOSPA estimate.
- If we take the average over all of the particles, we get

$$\mathbf{o}_M = \sum_{i=1}^{N_H} w_i \mathbf{o}_i$$

This is the set of probabilities telling us the ordering of the MMOSPA estimate that would be "closest" (in an OSPA sense) to a random sample of the *underlying* set of particles.

- Suppose that the ordering is known exactly (i.e. $\mathbf{o}_i$ contains a single one and all zeros.
- Then we can represent $\mathbf{o}_i$ as a $N_T \times N_T$ permutation matrix, χ_i (containing exactly one 1 in each row and column).
- As has been discussed in previous work, if we mix the permutation matrices, the result can be interpreted as a matrix of marginal probabilities. e.g

$$\chi_M = \begin{bmatrix} & \text{State 1} & \text{State 2} & \text{State 3} \\ \text{Track 1} & 2/3 & 1/6 & 1/6 \\ \text{Track 2} & 1/3 & 1/3 & 1/3 \\ \text{Track 3} & 0 & 1/2 & 1/2 \end{bmatrix}$$

so

$$\Pr\{\text{State 1 is Track 1}\} = 2/3$$

If we make an independence assumption, then we can say

$$\Pr\{\text{State 1 is Track 1} \cap \text{State 2 is Track 2} \cap \text{State 3 is Track 3}\} \approx \frac{\Pr\{\text{State 1 is Track 1}\}\Pr\{\text{State 2 is Track 2}\}\Pr\{\text{State 2 is Track 2}\}}{\text{perm}[\chi_M]}$$

This formulation makes mixing simpler than using the full $\mathbf{o}_i$.
Note that a label-free tracker with labeling is not the same as a labeled tracker. The difference is that target locations are estimated without a labeling constraint; and the labeling identity is an "overlay".
It is worth reiterating that traditional (labeled) trackers *do* have uncertain labels, but keep quiet about them and deliver estimates as if there were no uncertainty.

Simulation: Crossing & Recrossing Targets



Separation cases 1.5 unit (upper) and 1.0 unit (lower). Left: JPDAF. Middle: SJPDAF. Right: SJPDA identity probabilities.

Monte Carlo Performance

Cross ?	min. sep.	SJPDA % Track loss	SJPDA P_l	SJPDA $\bar{P}^L(1)$	JPDA % Track loss	JPDA P_l
No	0.5	0	48.20	50.00	72.70	72.89
	1.0	0.10	54.05	53.54	0	56.10
	1.5	0	97.00	93.15	0	86.20
Yes	0.5	0	46.80	50.00	97.30	25.93
	1.0	0.10	60.36	50.85	3.40	37.68
	1.5	0.10	83.38	59.74	0.10	93.89

P_l is probability that both tracks have the same identity at the end as at the beginning. $\bar{P}^L(1)$ is SJPDA's (averaged) estimated probability that that both tracks have the same identity at the end as at the beginning.

GMR-MHTs

- Consider Salmond's mixture reduction idea[13]:
  - what MHT programmers call a hypothesis a statistician would call a mode of a Gaussian mixture
  - idea is to find another lower-complexity Gaussian mixture that matches the original one in the sense of ...
  - ... integrated square error[14] between pdfs – Kullback-Leibler or Bhattacharyya distance would be better, but ISE is explicit for Gaussians
- Recently shown[15] that Runnall's algorithm is best
- Note that GMR-MHT reflects uncertainty caused by GMR by inflating the covariances. For a regular MHT this would work for merged hypotheses, but not for pruned ones.

[13] D. Salmond, "Mixture reduction algorithms for point and extended object tracking in clutter," *IEEE T-AES*, 2009.

[14] J. Williams and P. Maybeck, "Cost-function-based hypothesis control techniques for multiple hypothesis tracking," *Mathematics of Computer Modeling*, 2006

[15] D. Crouse, P. Willett, K. Pattipati and L. Svensson, "A look at Gaussian mixture reduction," *FUSION*, 2011?

The SMHT[16]

- As just seen, the MHT is good at avoiding coalescence
  - "Soft" association algorithms (PDA based) are prone to coalescence.
  - "Hard" association algorithms like the MHT have track repulsion.
- But it is natural to combine the ideas of label-free tracking and the MHT, via GMR-MHT:
  - expand hypotheses by considering all permutations of targets
  - ignore unlikely permutations
  - use GMR to reduce the hypotheses from this expanded set
- Identities can be tracked here as with the SJPDAF.

[16] D. Crouse, P. Willett, L. Svensson, D. Svensson and M. Guerriero, "The Set MHT," *FUSION*, 2011?

Simulation



The paths of the tracks in the three scenarios. In (a), the targets cross at the beginning; in (b) the targets never cross and in (c) the targets cross twice. All units are in meters. These scenarios are fairly easy to track compared to the previous ones.

More Details Than are Useful to Mention

- A single sensor taking polar measurements was placed at $(3.39\,\text{km}, -30\,\text{km})$.
- Polar measurements were taken with $\sigma_r = 2\,\text{m}$ and $\sigma_\theta = 1\,\text{mrad}$.
- A sampling rate of $T = 1/2\,\text{s}$ was used.
- Clutter was generated according to a Poisson process over the viewing area with a density of $\lambda = 1 \times 10^{-6} V$ where V is the area of the surveillance region.
- The viewing region was the extrema of the target trajectories plus $200\,\text{m}$ is each direction.
- There were an average of 18 clutter points per scan.
- All trackers gated measurements to a 99.97% confidence region.
- The turns were $5.3g$ coordinated turns.

Results

Algorithm	Scenario 1		Scenario 2		Scenario 3	
	% Track Loss	P_l	% Track Loss	P_l	% Track Loss	P_l
JPDAF	6.20	97.95	9.96	49.96	8.98	94.08
Set JPDAF	6.18	97.15	9.86	53.62	8.92	90.98
GMR MHT	7.38	61.89	8.82	47.97	10.18	64.68
Set MHT	0.94	98.72	1.34	76.16	1.10	96.92

- The Set MHT has the lowest track loss and the highest probability of a correct assignment in all scenarios.[17]
- The GMR MHT has a higher rate of track loss than the single hypothesis trackers.

[17] The assignment probabilities were calculated on runs where no method lost tracks.

Summary


Kalman
Bar-Shalom: PDAF
Reid: MHT
Mori: MHT
Blackman: MHT
Pattipati: SD Assignment
Streit & Luginbuhl: PMHT
Stone: UDF
Mahler: PHD
Mahler: CPHD
Salmond: Gaussian Sum Filter
You: RFS?
1960
1970
1980
1990
2000
2010
Lots of people: Explaining why Kalman Filtering is a Bad Idea
Morefield: LP solution
Bar-Shalom: JPDAF
Alphatech Track-Oriented MHT
JVC, Drummond: Efficient GNN
Poore: MF Assignment
Mahler: Random Sets
Kastella: JMP
Vo: GMPHD
Streit: IMPPP
Svensson: SJPDA